

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 695 062 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
31.01.1996 Bulletin 1996/05

(51) Int. Cl.⁶: H04L 12/56, H04Q 11/04

(21) Application number: 95111626.8

(22) Date of filing: 24.07.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 25.07.1994 US 280190

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)

(72) Inventors:
• Rashid, Richard F.
Woodinville, Washington 98072 (US)

• Bolosky, William J.
Issaquah, Washinton 98027 (US)
• Fitzgerald, Robert P.
Redmond, Washington 98053 (US)

(74) Representative: Grünecker, Kinkeldey,
Stockmair & Schwanhäusser
Anwaltssozietät
D-80538 München (DE)

(54) Method and system for combining data from multiple servers

(57) An ATM switch is used to combine blocks of data from multiple data servers into a single continuous data stream. The ATM switch is especially well adapted for use in a video-on-demand environment. Blocks of data may be striped across the data servers so that they are output from respective data servers in sequence. Transmission of data from the data servers to the ATM switch is scheduled so that the blocks of data are output in

sequence without undue delay between the transmission of blocks of data from the servers. Synchronizing tokens are provided to the data servers to specify that they are to output data of a given sequence and to direct the sequence to a given virtual channel. As a result, the system of this invention appears to output single continuous data streams from the output ports of the ATM switch.




FIG. 1

BEST AVAILABLE COPY

EP 0 695 062 A2

Description

Technical Field

This invention relates generally to data processing systems and, more particularly, to combining data from multiple data servers into a single continuous data stream using a switch.

Background of the Invention

Asynchronous transfer mode (ATM) is one of a general class of high-speed packet switching technologies that route traffic via an address contained within each packet. Asynchronous transfer mode may be contrasted with synchronous transfer mode techniques that are used in many digital telephone networks. Asynchronous transfer mode is an outgrowth of the broadband integrated services digital network (BISDN) standards that have been developed. ATM switches have been developed to control routing in ATM environments. These switches route fixed size packets, known as cells, to their intended destinations. Generally, ATM switches are used to realize a connection from a single origination point to a single destination point.

Summary of the Invention

In accordance with one aspect of the present invention, a method is practiced in a system that has a plurality of data servers for outputting data and an ATM switch. In this method, consecutive blocks of data in a data sequence are output from the data servers. The ATM switch is then used to combine the output blocks into a single sequential data stream.

In accordance with another aspect of the present invention, a first set of data is output from a first data server to an ATM switch. The first set of data is directed through the ATM switch to a selected output port on the ATM switch. A second set of data is output from a second data server to the ATM switch. The second set of data is also directed through the ATM switch to the selected port of the ATM switch to create a continuous stream of output with the first set of data from the selected output port.

In accordance with a further aspect of the present invention, a method is practiced in a video-on-demand server system having a plurality of data servers and a switch with input ports and output ports. In this method, a first block of data of a video image sequence is output from a first of the data servers to a first of the input ports of the switch. The first block of the video image sequence is directed through the switch to a first of the output ports. A second block of the video image sequence is output from a second of the data servers and directed through the switch to the first of the input ports to produce a continuous output data stream with the first block of the video image sequence.

In accordance with an additional aspect of the present invention, a token that identifies an output port

of an ATM switch to which cells are to be output is forwarded to a first data server. The cells holding data for a first block in a data sequence are output from the first data server to the output port of the ATM switch that is identified by the token. The token is then forwarded to a second of the data servers and the cells holding data for a second block of the data sequence are output from the second data server through the output port that is identified by the token.

In accordance with yet another aspect of the present invention, a method is practiced in a system that has a plurality of servers for outputting data to an asynchronous transfer mode (ATM) switch that directs data to virtual channels. In this method, a token that identifies a virtual channel is passed to a first of the data servers. In response to the token being received at the first of the data servers, an identifier of the virtual channel is incorporated in a header that is part of the cells of a first block of data of a data sequence. The cells for the first block of data are output from the first data server to the ATM switch so that the cells are routed over the identified virtual channel. A token that identifies the virtual channel is passed through a second of the data servers. In response to receiving the token at the second data server, the second data server incorporates the identifier of the virtual channel and the header for the cells of a second block of data of the data sequence. These cells in the second block of data then output from the second data server to the ATM switch so that they are routed over the identified virtual channel.

Brief Description of the Drawings

Figure 1 is a block diagram depicting a video-on-demand server system that is suitable for practicing a preferred embodiment of the present invention.

Figure 2 is a diagram depicting the timing of the outputting of blocks of data in accordance with the preferred embodiment of the present invention.

Figure 3 is a block diagram depicting the architecture of a conventional ATM switch.

Figure 4 is a diagram illustrating the notion of virtual paths and virtual channels in a conventional ATM environment.

Figure 5 is a block diagram illustrating the format of a conventional ATM cell.

Figure 6 is a flow chart of the steps performed to route a cell through a conventional ATM switch.

Figure 7 is a flow chart illustrating the steps performed to combine consecutive data blocks into a continuous data stream through the ATM switch in accordance with the preferred embodiment of the present invention.

Figures 8A and 8B illustrate an example of operation of an ATM switch to combine consecutive blocks of data into a single, continuous data stream in accordance with the preferred embodiment of the present invention.

Detailed Description of the Invention

The preferred embodiment of the present invention provides a method and system for combining data originating from multiple data servers into a single continuous data stream using an ATM switch. This combining of data may be concurrently performed on multiple data sequences to produce multiple continuous output data streams. Data transmissions from the data servers are synchronized for each data sequence through a combination of the use of synchronizing tokens and time division multiplexing of outputs from the data servers. The method and system of the preferred embodiment of the present invention allows a distributed data server system, such as a video-on-demand server system, that is built out of two or more computers, to appear to a client as a single computer sending a single data stream for each data sequence. The preferred embodiment greatly simplifies downstream switching of the data stream (i.e., switching after the data leaves the server system).

The preferred embodiment of the present invention will be described below as being implemented in a video-on-demand server system. Nevertheless, those skilled in the art will appreciate that the present invention may also be practiced in other environments wherein multiple data servers produce data for a data sequence that is to be combined into a single data stream. The preferred embodiment of the present invention will also be described below as being implemented using ATM switches. Those skilled in the art will appreciate that the present invention may also be practiced using other types of switching networks. For example, the present invention may also be practiced in cell-switched networks, circuit-switched networks, and non-ATM packet switched networks.

Figure 1 depicts a block diagram of a video-on-demand server system 10 that is suitable for practicing the preferred embodiment of the present invention. The video-on-demand server system 10 includes a controller 12 that controls the activities of a plurality of data servers 14. The output of the data servers is fed into an ATM switch 16. Each data server 14 is connected to a respective input port of the ATM switch 16. The ATM switch 16 combines data from the multiple servers 14 into data streams that are output to subscribers 18. Those skilled in the art will appreciate that in alternative embodiments the ATM switch 16 also handles input from at least one additional set of components containing a controller and servers 14. The connection between the ATM switch 16 and the subscribers 18 may include additional switching technology that is not depicted in Figure 1. Specifically, the ATM switching network need not connect directly to the subscribers; rather, different transport mechanisms, such as television cable, telephone switching networks, local area networks, or other interconnection mechanisms may connect the outputs of the ATM switch 16 to the subscriber 18.

Each server 14 includes at least one storage device for storing blocks of the video image sequences that are

available to the subscribers 18. The servers 14 may hold the video image sequences on storage devices such as magnetic disk drives, magnetic tape drives, optical disk drives, RAM, ROM, EPROM, and flash EPROMS. Each video image sequence is stored across the storage devices of the respective servers 14. Multiple copies of any video image sequence may be stored within the video-on-demand server system 10. The striping of the video image sequences increases the efficiency with which data is output by the storage devices in a fixed amount of time and helps to balance load requirements for each storage device. The consecutive blocks of each video image sequence are stored on the storage devices according to a predetermined sequence of storage devices. Nevertheless, those skilled in the art will appreciate that the present invention is not limited to embodiments wherein the data is striped across the servers 14. When striping is applied, an initial block of a video image sequence is stored on a first server, and the next block of the video image sequence is stored on a next storage device in a next one of the servers 14. The sequence of storing blocks of the video image sequence repeats itself until all the data for the video image sequence have been stored in a striping fashion across the storage devices. The previously determined sequence wraps around to the beginning when the end of the sequence is reached. The block size may vary but typically ranges between 64 K to 4 MB of data. Each block within a given data sequence has a same maximum size, but the maximum block size may vary between data sequences.

In the video-on-demand system, a subscriber 18 may request at any point in time to view a particular video image sequence that is transmitted from the servers 14 through the ATM switch 16 to the subscriber. The data for the video image sequences may contain video data, audio data and/or other types of data, including close-captioning data. The present invention may also be applied solely to video data, audio data or to other types of data sequences.

The display of video image sequences to the subscriber 18 is not pre-scheduled. Video image sequences are forwarded to the subscriber only upon demand. A subscriber need not choose a video image sequence that other subscribers are watching; rather, the subscriber may choose many of the available video image sequences that are provided by the system. Further, each subscriber chooses when he wishes to start viewing a video image sequence and may use controls, such as a pause control, a stop control, a fast forward control and a rewind control, to regulate his viewing of the video image sequence. A subscriber may select where in a sequence he desires to start viewing and stop viewing a video image sequence. A number of different subscribers may be concurrently viewing different portions of the same video image sequence. Alternatively, the video image sequences may be output not to just a single subscriber 18 but to multiple subscribers who cannot individually pause, rewind, stop or fast forward the video image sequence.

Figure 2 illustrates how blocks of data of video image sequences are output from the storage devices of the data servers 14. The scheduling of data transmission output from each of the disks is achieved using time division multiplexing. The scheduling is described in more detail in copending application No. 08/159,188, entitled "Method and System for Scheduling Transfer of Data Sequences," filed on November 11, 1993, and assigned to a common assignee with the present application. The outputting of blocks of data is divided according to time slots. Each time slot is a sufficient amount of time for a disk to output a block of data. Thus, as shown in Figure 2, disk 0 outputs a block of data in a first time slot 1 and outputs a second block of data in a second time slot. The numbers in the rectangles correspond with subscriber identifiers. Each subscriber receives a separate data stream of blocks of a video image sequence. The rectangles in Figure 2 represent the time slots. Each time slot is of sufficient length to output an entire block of data. The outputting of data by the disk (disk 0, disk 1 and disk 2) is staggered such that at the beginning of each time unit (a measure of a master clock), a different disk begins outputting a block of data. For example, at time zero, disk 1 begins outputting a block of data that is destined for a first subscriber. At time 1, disk 1 begins outputting the next block of the video data image sequence destined for subscriber 1. It should be recalled, as mentioned above, that the blocks of a video image sequence are striped across the disks. Accordingly, these three disks (disk 0, disk 1 and disk 2) may be concurrently outputting seven separate data streams to seven separate subscribers 18.

Those skilled in the art will appreciate that the present invention is not limited to environments in which a server is output data in a predefined sequence; rather, the present invention is well adapted for use in environments wherein multiple servers output data in a non-pre-determined sequence. Moreover, the present invention may also be used with non-time critical data that need not be output as a properly sequenced continuous data stream like the video data stream described above. Specifically, the present invention may be practiced in environments where the output from the data servers is not time division multiplexed.

In order to appreciate how the ATM switch 16 is used to combine data from multiple data servers 14 into a single output data stream for a data sequence, it is helpful to further examine the architecture of the ATM switch 16. Figure 3 shows the architecture of the ATM switch 16 in more detail. The ATM switch 16 includes a control processor 20, switching fabric 22, input controllers 24 and output controllers 26. In the example shown in Figure 3, the ATM switch 16 has seven input ports, seven input controllers 24, seven output controllers 26, and seven output ports. The control processor 20 is used to implement higher level functionality that is not the concern of the present invention. Each input controller 24 is associated with a corresponding input port. The input controllers 24 receive input from the corresponding input port

and manipulate the cells of information so that they are directed to the proper output port via the switching fabric 22. The input controllers 24 are responsible for examining the incoming cells that are received at the input ports and manipulating the cells so that they are properly routed to the output ports. The switching fabric 22 is comprised of a number of switching elements. The output controllers 26 control cells received from the switching fabric 22 and direct the cells to their corresponding output port.

It should be appreciated that the ATM switch configuration shown in Figure 3 is merely illustrative. Other ATM switch designs may also be used to practice the present invention. Those skilled in the art will also appreciate that a number of different types of ATM switches are known and may be used to practice the present invention.

In an ATM environment, routing is achieved by specifying virtual channels (VCs) and virtual paths (VPs) in cells. A "virtual channel" is a connection between two communicating ATM entities. "Virtual paths" are groups of virtual channels that are carried between two points. Virtual paths provide a convenient technique for bundling traffic that is heading for the same destination. In particular, switching equipment need only check for a virtual path identifier to relay traffic rather than checking an entire address contained within a cell. Figure 4 shows an example of the concepts of a virtual path and virtual channel. In the example shown in Figure 4, a physical link 28 carries traffic for three virtual paths: VP1, VP2, and VP3. Each of the virtual paths include one or more virtual channels. Specifically, VP1 includes virtual channels VC1, VC2, and VC3; VP2 includes virtual channels VC4 and VC5; and VP3 includes virtual channel VC6.

As mentioned above, in an ATM network environment, all communications are packaged into fixed size cells. Each cell holds information that identifies the virtual path and virtual channel over which the cell is to be routed. Each server 14 includes a network interface card that serves as an interface to the ATM network. The ATM interface provided on the network interface card packages the blocks of data that are retrieved from the storage devices of the servers 14 into cells that are suited to be transferred to the ATM switch 16. Figure 5 shows an example of an ATM cell. The ATM cell is 53 bytes long and includes a 5-byte header 32 and a 48-byte information field 34. The header includes a virtual path identifier (VPI) 36 and a virtual channel identifier (VCI) 38. The VPI 36 and VCI 38 are used by input controllers 24 (Figure 3) in the ATM switch 16 (Figure 1) to properly route cells to the output ports associated with the specified virtual paths and virtual circuits. It should be appreciated that the values of VPI and VCI stored within a cell may change as a cell is routed over an ATM network, as will be described in more detail below.

Figure 6 is a flow chart of the steps performed in routing a cell through an ATM switch. Initially, the incoming cell is received by the ATM switch at one of its input ports (step 40). The VPI and VCI stored in the header

32 (Figure 5) of the cell 30 are examined by the input controller 24 (Figure 3) of the ATM switch (step 42 in Figure 6). Each ATM switch 16 (Figure 1) contains a connection table, which is a lookup table used to translate the set containing the input port end VPI 36 and VCI 38 in the header of a cell to a new set that includes an output port, output VPI and output VCI so that the cell may be routed over the proper path to reach its destination. The connection table stores output port identifiers that are indexed by the VPI/VCI of an incoming cell. As such, the input controller 24 (Figure 3) accesses the connection table to determine the output port to which the cell should be destined (step 44 in Figure 6). The connection table may take the form of a global table that is shared by the input controllers 24 or alternatively each input controller may have its own table. The input controller 24 attaches a routing tag to the cell that is used by the switching fabric 22 to determine how to route the cell to the desired output port and overwrites the VPI 36 (Figure 5) and VCI 38 to values that are useful in the next ATM link, if any (step 46 in Figure 6). For purposes of simplification in the present example, it is assumed that the ATM connection includes only a single link.

The cell is then passed to the switching fabric 22 (step 48 in Figure 6) so that it may be routed to the proper output port (step 50). The switching fabric 22 may include buffers for buffering cells while they are being transmitted across the switching fabric.

In the preferred embodiment of the present invention, the blocks of data for a video image sequence that are striped across the servers 14 are merged into a single continuous data stream in the ATM switch 16 by assigning VPI/VCI values to realize such combining of the blocks of data into a single data stream and by using synchronizing tokens as described below. The cells of data for each block of data in a video image sequence are each assigned a common virtual path and virtual channel. Hence, the blocks of data for a video image sequence are directed to a same destination by the ATM switch 16. The synchronizing tokens ensure proper sequencing for cells that are output from different servers to a common virtual channel. In other words, the synchronizing tokens ensure that the cells are sent in proper order.

Figure 7 is a flow chart illustrating the steps that are performed by a server 14 to output a block of data for a video image sequence. Initially a server 14 is waiting for a synchronizing token so that it can output data (step 52). The synchronizing token is then passed to the server (step 54). The server 14 emits the enabled data block as cells through its network interface card to the ATM switch 16 (step 56). Those skilled in the art will appreciate that the token may arrive before the data is ready to be output from the server or vice versa. The flow chart of Figure 7 is merely illustrative in showing the data arriving first. The synchronizing token is then passed to the next server as the predetermined sequence of servers (i.e., the server that holds the next block of data of the video image server 10) so that the next server may emit data (step 58). Alter-

natively, the synchronizing token may be returned to a known location where the next server 14 will know to look. The steps of Figure 7 are independently repeated for each of the servers 14.

A separate token is provided for each virtual channel (i.e., each output data stream). This token is passed in sequence from server 14 to server so that the blocks of video image sequence are output in proper order. The synchronizing tokens are generated by a separate serving process that may reside on one of the servers 14 or on a separate computer. This serving process is responsible for generating, managing, and deleting the tokens for each of the virtual channels that is in use.

The use of the synchronizing tokens preserves data order in the presence of variable queuing delays such as found in multistage interconnection networks. In addition, the use of the synchronizing tokens allows the present invention to work in instances where several users were on at different data rates and with different blocking factors. For example, the present invention is well-adapted to be used in an environment wherein a different number of bytes are output from the respective servers 14.

Figures 8A and 8B illustrate an example of the operation of the preferred embodiment of the present invention in combining data for multiple data servers into a single data stream. These steps are separately performed for each output data stream. In the example shown in Figures 8A and 8B, each server has an associated input port. In Figure 8A, the server associated with input port 1 forwards cells for a first block of a video image sequence to the ATM switch 16. The cells are provided with VPI/VCI values that direct the ATM switch to forward the cells to the output port 1. During a next time frame, after all of the data for the first block of the video image sequence has been output, a token is passed to the second server that is associated with input port 2. The second server passes data via input port 2 into the ATM switch 16 and sets a switching logic so that the cells of the second block of the video image sequence are directed out output port 1. This process is repeated for the remaining servers that hold consecutive blocks of the video image sequence until the last server is reached at which point the sequence repeats itself.

While the present invention has been described with reference to a preferred embodiment thereof those skilled in the art will appreciate the various changes in form and detail may be made without departing from the intended scope of the claims as defined in the appended claims. For example, the present invention may be implemented with different switch architectures than those described above.

Claims

1. In a system having a plurality of data servers for outputting data and a virtual circuit network switch, a method comprising the steps of:
outputting consecutive blocks of data in a

data sequence from respective ones of the data servers; and

combining the blocks of data that are output by the data servers with the virtual circuit network switch into a single sequential data stream.

2. The method of claim 1 wherein the data servers include storage for storing the consecutive blocks of data in the data sequence and wherein the step of outputting consecutive blocks of data in sequence from respective ones of the data servers comprises the step of reading the consecutive blocks of data in the data sequence from the storage of the data servers. 10
3. The method of claim 1, further comprising the step of striping the blocks of data in the data sequence across the data servers so that consecutive blocks of data in the data sequence are stored on different data servers. 15 20
4. The method of claim 3 wherein the step of outputting consecutive blocks of data in the data sequence from respective ones of the data servers includes the step of time multiplexing the outputting of the consecutive blocks of data in the data sequence so that each data server is given a time slot in which it outputs a block of the data sequence stored therein. 25
5. The method of claim 1 wherein the virtual circuit network switch is an asynchronous transfer mode (ATM) switch. 30
6. In a system having data servers and a virtual circuit network switch having output ports, a method comprising the steps of: 35
 - outputting a first set of data from a first of the data servers to the virtual circuit network switch;
 - directing the first set of data through the virtual circuit network switch to a selected output port of the virtual circuit network switch; 40
 - outputting a second set of data from a second of the data servers to the virtual circuit network switch; and
 - directing the second set of data through the virtual circuit network switch to the selected port of the virtual circuit network switch so as to create a continuous stream of output with the first set of data from the selected output port. 45
7. The method of claim 6 wherein the first set of data and the second set of data are part of a same data sequence. 50
8. The method of claim 6 wherein the step of outputting the first set of data from the first of the data servers to the virtual circuit network switch comprises the step of outputting a first set of cells holding the first set of data and information that identifies a virtual 55

channel to which the first set of cells are destined from the first of the data servers.

9. The method of claim 8 wherein the step of directing the first set of data through the virtual circuit network switch to the selected output port of the virtual circuit network switch includes the steps of:
 - examining the information in the first set of cells that identifies the virtual channel to which the cells are destined to identify the virtual channel to which the first set of cells are destined;
 - determining that the selected output port is associated with the virtual channel to which the first set of cells are destined; and
 - directing the first set of cells through the virtual circuit network switch to the selected output port.
10. The method of claim 6 wherein the step of outputting the second set of data from the second of the data servers to the virtual circuit network switch comprises the step of outputting a second set of cells holding the second set of data and information that identifies a virtual channel to which the second set of cells are destined from the second of the data servers.
11. The method of claim 10 wherein the step of directing the second set of data through the virtual circuit network switch to the selected output port of the virtual circuit network switch includes the steps of:
 - examining the information in the second set of cells that identifies the virtual channel to which the cells are destined to identify the virtual channel to which the second set of cells are destined;
 - determining that the selected output port is associated with the virtual channel to which the second set of cells are destined; and
 - directing the second set of cells through the virtual circuit network switch to the selected output port.
12. The method of claim 6 wherein the virtual circuit network switch is an asynchronous transfer mode (ATM) switch.
13. In a video-on-demand server system having a plurality of data servers and a switch having input ports and output ports, a method comprising the steps of:
 - outputting a first block of a video image sequence from a first of the data servers to a first of the input ports of the switch, wherein the video image sequence comprises a sequence of blocks of data;
 - directing the first block of the video image sequence through the switch to a first of the output ports of the switch; and
 - outputting a second block of the video image sequence from a second of the data servers;

directing the second block of video image sequence through the switch to the first of the input ports of the switch to produce a continuous output data stream with the first block of the video image sequence from the first of the output ports.

5

14. The method of claim 13 wherein the video image sequence includes audio data and video data.

15. The method of claim 13, further comprising the steps of:

10

outputting additional blocks of the video image sequence from the data servers;

directing the additional blocks of the video image sequence through the switch to the first of the output ports to produce a continuous output stream of the blocks of the video image sequence.

15

16. The method of claim 13 wherein the switch is an asynchronous transfer mode switch.

20

17. In a computer system having data servers and a switch with output ports, a method of outputting a sequence of blocks of data comprising the steps of:

forwarding a token to a first of the data servers, said token identifying an output port of the switch to which cells are to be output;

25

in response to receiving the token, outputting cells holding data for a first block in the sequence from the first of the data servers to the output port of the switch that is identified by the token;

30

forwarding the token to a second of the data servers; and

in response to receiving the token, outputting cells holding data for a second block in the sequence from the second of the data servers to the output port of the switch that is identified by the token.

35

18. The method of claim 17 wherein the switch is an asynchronous transfer mode switch.

40

19. In a system having a plurality of servers for outputting data to an asynchronous transfer mode (ATM) switch that directs data to virtual channels, a method comprising the steps of:

45

passing a token that identifies a virtual channel to a first of the data servers;

in response to receiving the token at the first of the data servers, incorporating an identifier of the virtual channel in a header that is incorporated into cells of a first block of data of a data sequence;

50

outputting the cells for the first block of data from the first of the data servers to the ATM switch so that the cells for the first block of data are routed over the identified virtual channel;

55

passing the token that identifies the virtual channel to a second of the data servers;

in response to receiving the token at the second of the data servers, incorporating the identifier

of the virtual channel in a header that is incorporated into cells of a second block of data of the data sequence; and

outputting the cells from the second block of data from the second of the data servers to the ATM switch so that cells for the second block of data are routed over the identified virtual channel.




FIG. 1




FIG. 2

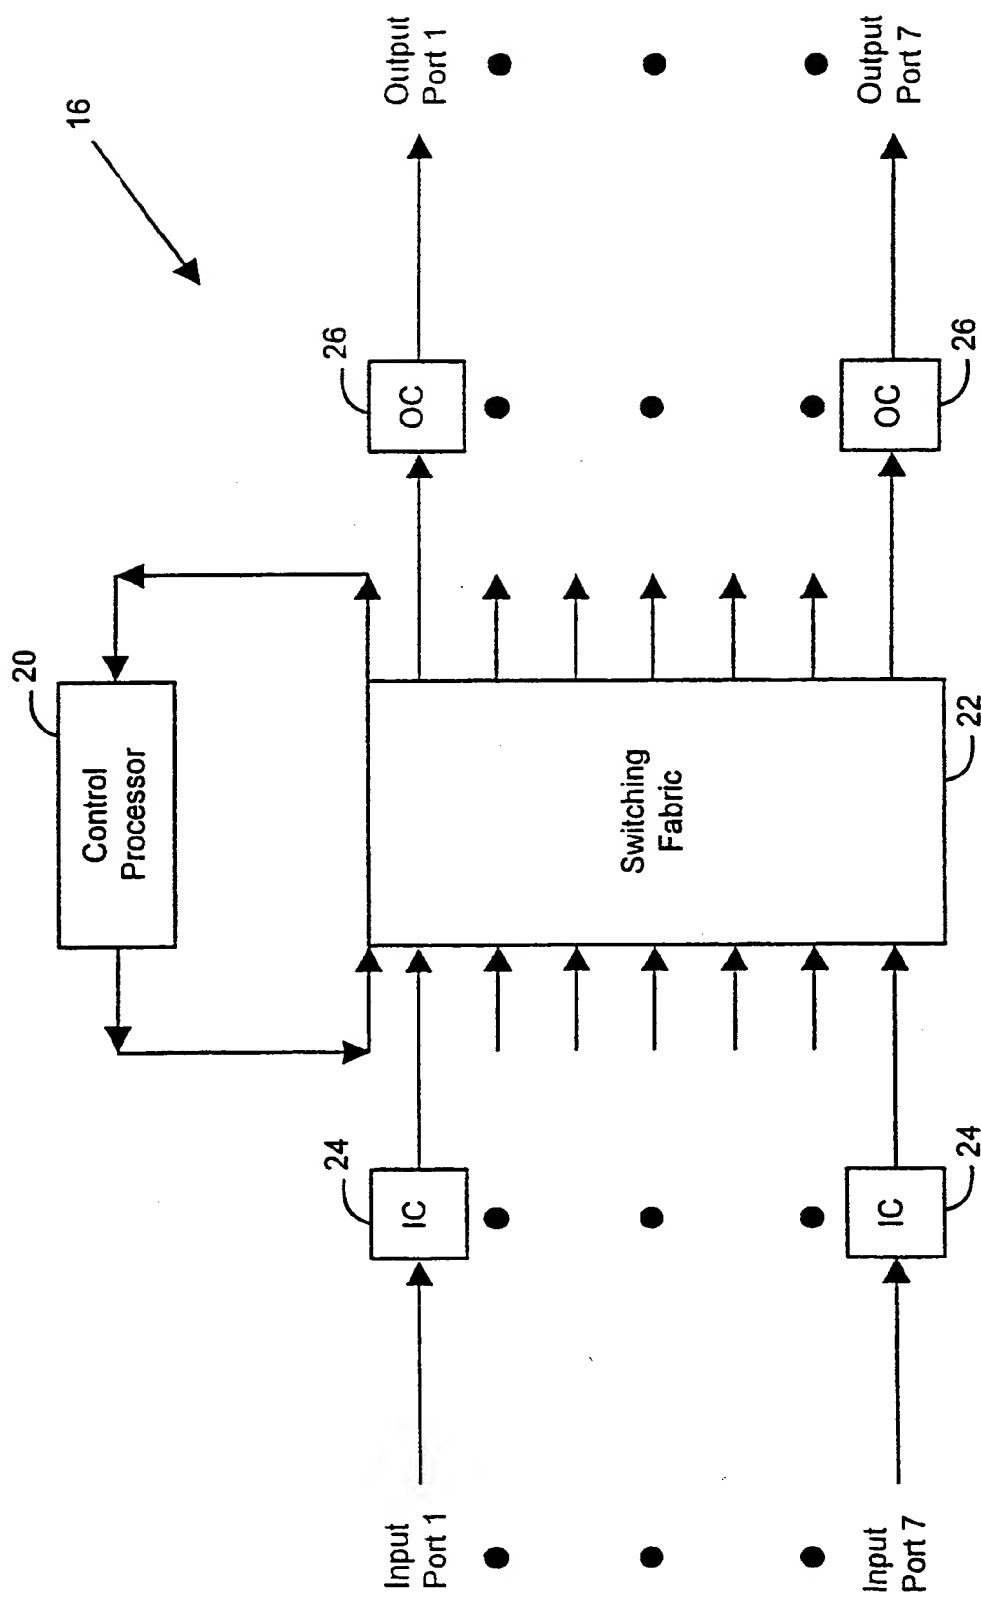




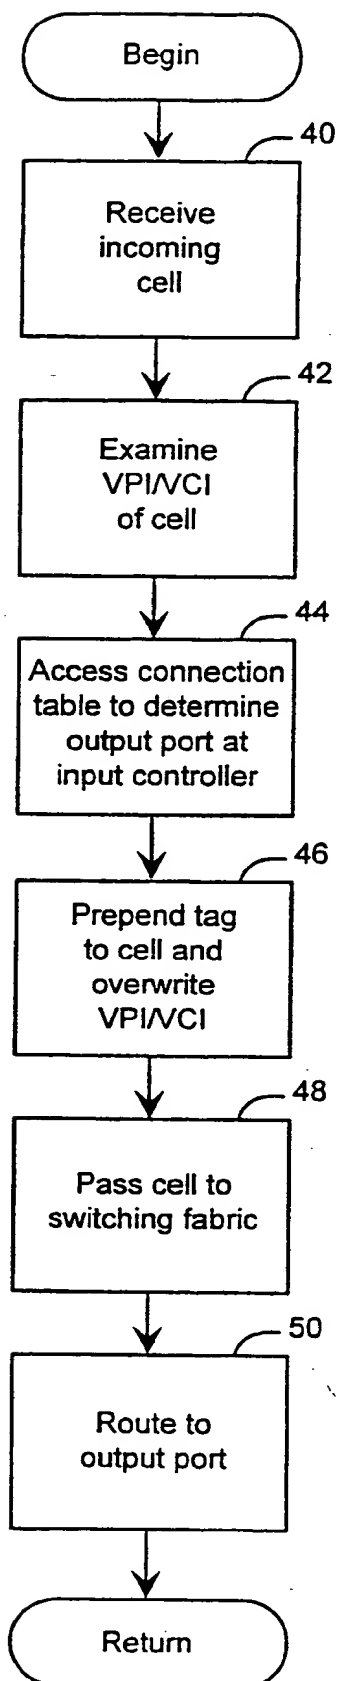
FIG. 3



PRIOR ART
FIG. 4



PRIOR ART
FIG. 5

**FIG. 6**

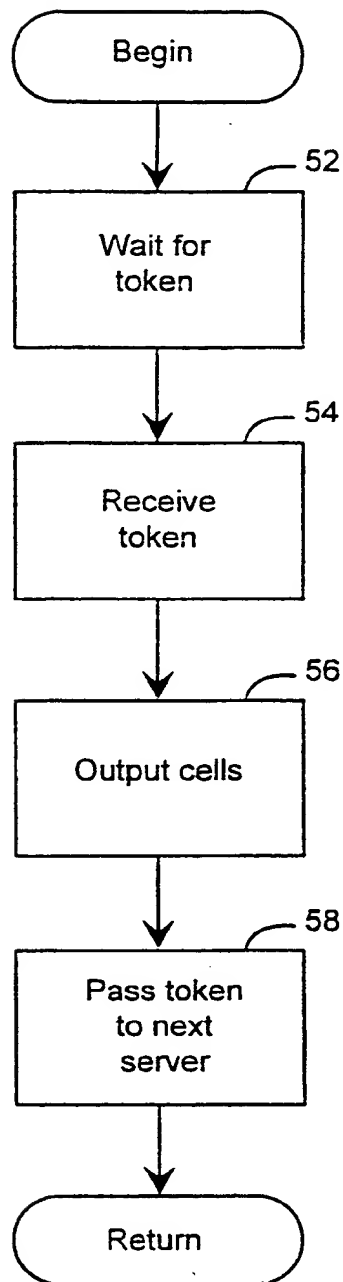


FIG. 7




FIG. 8A





FIG. 8B

THIS PAGE BLANK (USPTO)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 695 062 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
21.06.2000 Bulletin 2000/25

(51) Int. Cl.⁷: H04L 12/56, H04Q 11/04

(43) Date of publication A2:
31.01.1996 Bulletin 1996/05

(21) Application number: 95111626.8

(22) Date of filing: 24.07.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 25.07.1994 US 280190

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)

(72) Inventors:
• Rashid, Richard F.
Woodinville, Washington 98072 (US)

• Bolosky, William J.
Issaquah, Washinton 98027 (US)
• Fitzgerald, Robert P.
Redmond, Washington 98053 (US)

(74) Representative:
Grünecker, Kinkeldey,
Stockmair & Schwanhäusser
Anwaltssozietät
Maximilianstrasse 58
80538 München (DE)

(54) Method and system for combining data from multiple servers

(57) An ATM switch is used to combine blocks of data from multiple data servers into a single continuous data stream. The ATM switch is especially well adapted for use in a video-on-demand environment. Blocks of data may be striped across the data servers so that they are output from respective data servers in sequence. Transmission of data from the data servers to the ATM switch is scheduled so that the blocks of data are output in sequence without undue delay between the transmis-

sion of blocks of data from the servers. Synchronizing tokens are provided to the data servers to specify that they are to output data of a given sequence and to direct the sequence to a given virtual channel. As a result, the system of this invention appears to output single continuous data streams from the output ports of the ATM switch.




FIG. 1

EP 0 695 062 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 1626

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	ITO YUKIKO, TANAKA TSUTOMU: "A VIDEO SERVER USING ATM SWITCHING TECHNOLOGY" MULTIMEDIA' 94. 5TH IEEE COMSOC INTERNATIONAL WORKSHOP ON MULTIMEDIA COMMUNICATIONS, 16 May 1994 (1994-05-16), pages 3/4/1-3/4/6, XP002113965 retrieved from internet: <URL: http://iel.ihs.com:80/cgi-bin/iel.cgi >	1-3,5-8, 10,12-16	H04L12/56 H04Q11/04
Y	* page 3-4-1, column 1, line 7 - line 14 * * page 3-4-1, column 2, line 16 - line 21 * * page 3-4-1, column 2, line 39 - page 3-4-2, column 1, line 9 * * page 3-4-2, column 1, line 14 - line 20 * * page 3-4-2, column 2, paragraph 3.1.1 - page 3-4-3, column 1 * * page 3-4-4, column 2, line 24 - line 32 * * figures 1,2 *	4,9,11, 17,18	
Y	SCHNEIDER H: "THE CONCEPT OF VIRTUAL PATHS AND VIRTUAL CHANNELS IN ATM-NETWORKS" ELECTRONIC CIRCUITS AND SYSTEMS FOR COMMUNICATION, ETH ZURICH, MAR. 5 - 8, 1990, 5 March 1990 (1990-03-05), pages 63-72, XP000315092	9,11	
A	MELCHIOR H * abstract * * page 63, column 2, paragraph 2. - page 64, column 1 * * page 65, column 2, line 43 - page 66, column 1, line 10 * * figure 2.1 * --- -/--	8,19	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 3 April 2000	Examiner Lamadie, S
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document</p> <p>T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document</p>			

EPO FORM 1503 03/82 (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 1626

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
Y	TOSHIYA ARAMAKI ET AL: "PARALLEL "ATOM" SWITCH ARCHITECTURE FOR HIGH-SPEED ATM NETWORKS" DISCOVERING A NEW WORLD OF COMMUNICATIONS, CHICAGO, JUNE 14 - 18, 1992, vol. 1, 14 June 1992 (1992-06-14), pages 250-254, XP000326879 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS ISBN: 0-7803-0599-X * figures 1,3,5 * * page 311.1.2, column 1, line 6 - line 12 * * page 311.1.2, column 2, paragraph 3.2 *	4	
Y	SCHULTZ K J ET AL: "DISTRIBUTED MULTICAST CONTENTION RESOLUTION USING CONTENT ADDRESSABLE FIFOs" INTERNATIONAL CONFERENCE ON COMMUNICATIONS (ICC), US, NEW YORK, IEEE, vol. -, 1994, pages 1495-1500, XP000438745 * figure 1 * * page 1495, column 2, line 26 - page 1496, column 1, line 7 *	17,18	
A	---	19	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 3 April 2000	Examiner Lamadie, S
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03/82 (P04/001)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 1626

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.CI.6)
A	<p>LOUGHER P ET AL: "THE DESIGN OF A STORAGE SERVER FOR CONTINUOUS MEDIA"</p> <p>COMPUTER JOURNAL, vol. 36, no. 1, 1 January 1993 (1993-01-01), pages 32-42, XP000360263 ISSN: 0010-4620</p> <p>* figures 1-4 *</p> <p>* page 33, column 1, line 45 - line 51 *</p> <p>* page 33, column 2, line 3 - line 6 *</p> <p>* page 33, column 2, line 26 - line 34 *</p> <p>* page 34, column 2, line 11 - line 22 *</p> <p>* page 37, column 2, paragraph 5. - page 38, column 1 *</p> <p>* page 38, column 1, line 30 - line 34 *</p> <p>* page 38, column 2, line 4 - line 16 *</p> <p>* page 39, column 1, line 13 - line 39 *</p> <p>* page 39, column 1, line 46 - line 52 *</p> <p>* page 41, column 1, line 4 - line 7 *</p>	2,3	
A	<p>DELODDERE D ET AL: "INTERACTIVE VIDEO ON DEMAND"</p> <p>IEEE COMMUNICATIONS MAGAZINE, vol. 32, no. 5, 1 May 1994 (1994-05-01), pages 82-88, XP000451098 ISSN: 0163-6804</p> <p>* figures 1,3,6 *</p> <p>* page 82, column 2, line 13 - line 36 *</p> <p>* page 83, column 1, line 13 - line 16 *</p> <p>* page 83, column 1, line 43 - column 2, line 12 *</p> <p>* page 84, column 2, line 21 - line 33 *</p> <p>* page 85, column 2, line 1 - line 18 *</p>	1,5,6, 12,13,16	<p>TECHNICAL FIELDS SEARCHED (Int.CI.6)</p>
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
THE HAGUE		3 April 2000	Lamadie, S
CATEGORY OF CITED DOCUMENTS		<p>T : theory or principle underlying the invention</p> <p>E : earlier patent document, but published on, or after the filing date</p> <p>D : document cited in the application</p> <p>L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>	
<p>X : particularly relevant if taken alone</p> <p>Y : particularly relevant if combined with another document of the same category</p> <p>A : technological background</p> <p>O : non-written disclosure</p> <p>P : intermediate document</p>			

EPO FORM 1503 03 92 (P04001)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 1626

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	<p>EP 0 571 153 A (XEROX CORP) 24 November 1993 (1993-11-24) * figures 1,2 * * column 4, line 14 - line 21 * * column 5, line 11 - line 19 * * column 6, line 38 - line 50 * * column 8, line 38 - line 40 * * column 9, line 1 - line 3 * * column 9, line 28 - line 45 *</p>	19	
A	<p>US 5 276 681 A (TOBAGI FOUAD A ET AL) 4 January 1994 (1994-01-04) * column 3, line 32 - line 37 * * column 3, line 53 - line 65 * * column 4, line 1 - line 3 * * column 4, line 27 - line 47 * * column 5, line 33 - line 51 *</p>	17, 19	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 3 April 2000	Examiner Lamadie, S
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1603 03 92 (P04/C01)



European Patent
Office

Application Number

EP 95 11 1626

CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

- ☐ Only part of the claims have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid, namely claim(s):
- ☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

see sheet B

- ☒ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.
- ☐ As all searchable claims could be searched without effort justifying an additional fee, the Search Division did not invite payment of any additional fee.
- ☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid, namely claims:
- ☐ None of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims, namely claims:



European Patent
Office

**LACK OF UNITY OF INVENTION
SHEET B**

Application Number
EP 95 11 1626

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

1. Claims: 1-16

A method for outputting blocks of data from respective ones of the data servers.

2. Claims: 17-19

A method of outputting a sequence of blocks from servers connected to a switch by using a token identifying an output port of said switch.

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 95 11 1626

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

03-04-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0571153 A	24-11-1993	US 5327420 A	05-07-1994
		JP 6037800 A	10-02-1994
US 5276681 A	04-01-1994	NONE	

EPO FORM P0499

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)